Fitting Mixed Models in SPSS

You will find some background here on fitting a Linear Mixed Model in SPSS.

These models are useful for hierarchical data structures, in which the observations come in "groups". There are many ways in which grouping can occur, but for the purpose of presentation here, I will write about settings where observations are grouped by subject, with several observations on each subject.

The statistical model for the data can involve explanatory variables at the subject level (between subjects) or at the individual observation level (within subjects). A useful way to specify the models is to begin with a model for the results on a particular subject (the Level 1 model) and then to state how they relate to the between subject variables (the Level 2 model).

As an example, suppose that FEV_{ij} is the forced expiration volume from the i^{th} subject on the j^{th} testing occasion. There is one between subject factor, Ast, which is an indicator of whether the subject has asthma; and there is one within subject factor, PM, which is a measure of air pollution, the small particulate concentration in the atmosphere on the day of testing. At Level 1 a natural choice is a regression model relates FEV for the i^{th} subject to PM:

$$FEV_{ij} = \beta_{0i} + \beta_{1i} PM_{ij} + \varepsilon_{ij}.$$

Here I have suggested a model in which both the intercept and the slope are subject-specific.

Then the Level 2 model relates the results to the subject level variables, in this case whether or not the subject has asthma. This part of the model can include any features of the Level 1 model that are subject-specific. So here we would want models for both the intercept and the slope.

$$\beta_{0i} = \beta_0 + \beta_{0A} AST_i + \delta_{0i}.$$

$$\beta_{1i} = \beta_1 + \beta_{1A} AST_i + \delta_{1i}.$$

The terms δ_{0i} and δ_{1i} are error terms with a mean of 0. So the first Level 2 equation says that the intercepts have a typical value of β_0 for subjects who do not have asthma and a typical value of $\beta_0 + \beta_{0A}$ for those with asthma, and that they vary from one subject to the next, with δ_{0i} describing the inter-subject variability. So if you could directly measure the intercept for each non-asthmatic subject, and you did so for a sample of subjects, a histogram of the intercepts would be centered at β_0. The analysis will focus on estimating the features of these distributions, namely the central values β_0 and $\beta_0 + \beta_{0A}$, along with the spread of the distribution, as reflected in its variance σ_0^2. The equation for the slope has the same structure and reflects the fact that we want to take into account that not all subjects will have the same sensitivity to PM, even after we explain differences in sensitivity between asthmatic and non-asthmatic subjects.

It is perfectly legitimate to substitute the Level 2 equation into the Level 1 equation to get a combined equation that reads:

$$FEV_{ij} = \beta_0 + \beta_{0A} AST_i + \beta_1 PM_{ij} + \beta_{1A} PM_{ij} AST_i + \delta_{0i} + \delta_{1i} PM_{ij} + \varepsilon_{ij}.$$

We get a model that has a "main effect" for the Asthma indicator, a "main effect" for the level of pollution on the day of the test, an "interaction" between Asthma and PM, and three different error terms, two at the subject level and one at the single observation level.

Do we really need all these terms?? No, not always. For example, maybe asthma subjects and non-asthmatics are equally sensitive to air pollution, so there is no need for the interaction term. Maybe there is no inter-subject variability at all in sensitivity, so we don't need the subject-level error term associated with the slope on PM. These questions can be addressed by fitting the larger model and then assessing the output.

To fit models like the one above in SPSS, select Mixed Models and then Linear from the Analysis menu.

The first screen asks for Subject (the variable that identifies which outcomes are in the same group) and Repeated (the variable that numbers the observations within a particular subject).

It also prompts for the Repeated Covariance Type. This refers to how the error terms in the "Level 1" model (the ε_{ij} terms that we wrote in the model) are correlated with one another; it also lets you specify if they have the same variance or might have different variances. There are many options here and I will note just a couple.

Diagonal – the errors are independent and each one might have a different variance. No reason to think that the variance of FEV will change from one occasion to the next, but some outcome variables do show different variances.

Compound Symmetry – the errors all have the same variance and each pair of observations is correlated, with the same correlation. This is the error model you would get if you took the subject level error term for the intercept, δ_{0i}, and wrote it as part of the Level 1 model instead of making it a separate Level 2 model. The common "subject specific intercept" is what leads to having the same correlation between each pair of observations on the same subject.

AR(1) – the errors have the same variance and they are correlated, with the correlation high for "nearby" observations and increasingly smaller for more distant observations. This is sometimes used for longitudinal data, with the idea that errors from observations at nearby time points will be similar, but those from distant time points will not be similar.

Scaled Identity – the error terms have the same variance and they are not correlated.

Click continue to move to the next screen.

You are now prompted for the dependent variable, any factors (i.e. explanatory variables that are categorical), covariates (numerical explanatory variables) and residual weight (which I admit I do not understand). You should enter here any variables that you want to have in your equations, whether they are Level 1 or Level 2 equations.

Click on "Fixed" to set up the model terms you want in the regression equation. In the above example, we would have AST and PM as the two variables and we would want a model with both main effects and the interaction. We can specify this by choosing the "factorial" option when we enter AST and PM into the model. Other options are available. For example, if we chose "main effects" we would get a model that does not have the interaction term.

Return to the previous screen and click on "Random" to define the random effects that you want. You need to declare here each term *except* for the ε_{ij}, which are included automatically. In the above example, you would select PM as a random effect and you would check the "include intercept" box to indicate that the intercept varies from subject to subject. (As noted above, you could, instead, leave the box unchecked but specify "compound symmetry" as the variance structure for the Level 1 model; this gets the same results.)

At the bottom of the screen you are asked to specify the "subject groupings" and would need to move the "subject" variable to the right side of the window.

Click back to the earlier menu. Select "Statistics" to specify what output you want. I advise getting parameter estimates and tests for covariance parameters. You might also want descriptive statistics, though that output can be long.

Click on "EM Means" to specify comparisons that you would like to make between groups in the data. This will compare mean values for the groups, adjusting for other variables in the model, and also taking account of the complex error structure that comes from the hierarchical nature of the data. You can also choose from a couple different methods for accounting for multiple comparisons.

Example

The mixed model analysis is illustrated here by an experiment to assess the effects of two different drugs (A and B) on the amount of a particular substance in the blood. There were 4 experimental conditions: neither drug used, only A, only B or both A and B. The amount of the substance is noted in the file as the variable Y.

Eight mice were observed under each of the four conditions. The mice are the experimental units here and the data are "nested" within mouse. The "within mouse" data are arranged as a 2 by 2 factorial experiment, with one factor "use of drug A" and the other factor "use of drug B".

Four mice were males and four were females, so the analysis will also consider possible differences related to the sex of the mouse.

For the hierarchical experiment that we analyzed using Repeated Measures ANOVA, I chose to include all main effects of Sex, A and B (no interactions) and to include an error term only for the overall subject level but not for any of the effects. For the Level 1 model, I chose the "scaled identity" option. I also asked to compare means for A, B and Sex. The code looks like this:

```
MIXED Y BY Sex A B
  /CRITERIA=CIN(95) MXITER(100) MXSTEP(5) SCORING(1)
SINGULAR(0.000000000001) HCONVERGE(0,
    ABSOLUTE) LCONVERGE(0, ABSOLUTE) PCONVERGE(0.000001,
ABSOLUTE)
  /FIXED=Sex A B | SSTYPE(3)
  /METHOD=REML
  /PRINT=SOLUTION TESTCOV
  /RANDOM=INTERCEPT  | SUBJECT(Mouse) COVTYPE(VC)
  /REPEATED=Repeat | SUBJECT(Mouse) COVTYPE(ID)
  /EMMEANS=TABLES(A) COMPARE ADJ(LSD)
  /EMMEANS=TABLES(B) COMPARE ADJ(LSD)
  /EMMEANS=TABLES(Sex) COMPARE ADJ(LSD).
```

Here are the parameter estimates from the model. The model is set up so that the cell with Males, A=1, B=1 is the "reference" cell. The intercept of 6.06 gives the "best fit value" for that cell, assuming (as we have done here) that there are no interactions present. The estimate for A of -1.22 indicates that, for this model, *not* giving drug A, on average, leads to results that are 1.22 units lower than those achieved when A is administered. Similarly, *not* giving drug B leads to results that are 1.14 units lower. And females, on average, have 1.17 units more than males. The model recognizes that sex is at the "between subjects" level and assigns it 6 degrees of freedom and a larger SE. The effects for A and B each have smaller SE's (because they neutralize the between subject variation). The 22 degrees of freedom reflect 7 each from the "within animal" effects of A and B plus all 8 associated with the "within animal" A*B interaction (since I did not include that term in the model).

Estimates of Fixed Effects[b]

Parameter	Estimate	Std. Error	df	t	Sig.	95% Confidence Interval	
						Lower Bound	Upper Bound
Intercept	6.062500	.703161	9.985	8.622	.000	4.495443	7.629557
[Sex=Female]	1.168750	.870008	6	1.343	.228	-.960083	3.297583
[Sex=Male]	0[a]	0	.	.	.	.	.
[A=0]	-1.218750	.481619	22	-2.531	.019	-2.217566	-.219934
[A=1]	0[a]	0	.	.	.	.	.
[B=0]	-1.143750	.481619	22	-2.375	.027	-2.142566	-.144934
[B=1]	0[a]	0	.	.	.	.	.

a. This parameter is set to zero because it is redundant.

b. Dependent Variable: Y.

The estimated variance components are given in the following table. The top line, with an estimate of 1.86, corresponds to the variation at the within subject level. The next line, with an estimate of 1.05, corresponds to the variance, across animals, in the mean level of Y.

Estimates of Covariance Parameters[a]

Parameter		Estimate	Std. Error	Wald Z	Sig.	95% Confidence Interval	
						Lower Bound	Upper Bound
Repeated Measures	Variance	1.855653	.559501	3.317	.001	1.027661	3.35076
Intercept [subject = Mouse]	Variance	1.049915	.885131	1.186	.236	.201163	5.47973

a. Dependent Variable: Y.

The following tables compare the levels of A, then B and then Sex.

Estimates[a]

A	Mean	Std. Error	df	95% Confidence Interval	
				Lower Bound	Upper Bound
0	4.856	.497	9.985	3.748	5.964
1	6.075	.497	9.985	4.967	7.183

Estimates[a]

A	Mean	Std. Error	df	95% Confidence Interval	
				Lower Bound	Upper Bound
0	4.856	.497	9.985	3.748	5.964
1	6.075	.497	9.985	4.967	7.183

a. Dependent Variable: Y.

Pairwise Comparisons[b]

(I) A	(J) A	Mean Difference (I-J)	Std. Error	df	Sig.[a]	95% Confidence Interval for Difference[a]	
						Lower Bound	Upper Bound
0	1	-1.219*	.482	22	.019	-2.218	-.220
1	0	1.219*	.482	22	.019	.220	2.218

Based on estimated marginal means

*. The mean difference is significant at the .05 level.

a. Adjustment for multiple comparisons: Least Significant Difference (equivalent to no adjustments).

b. Dependent Variable: Y.

Estimates[a]

B	Mean	Std. Error	df	95% Confidence Interval	
				Lower Bound	Upper Bound
0	4.894	.497	9.985	3.786	6.002
1	6.038	.497	9.985	4.929	7.146

a. Dependent Variable: Y.

Pairwise Comparisons[b]

(I) B	(J) B	Mean Difference (I-J)	Std. Error	df	Sig.[a]	95% Confidence Interval for Difference[a]	
						Lower Bound	Upper Bound
0	1	-1.144*	.482	22	.027	-2.143	-.145
1	0	1.144*	.482	22	.027	.145	2.143

Based on estimated marginal means

*. The mean difference is significant at the .05 level.

a. Adjustment for multiple comparisons: Least Significant Difference (equivalent to no adjustments).

b. Dependent Variable: Y.

Estimates[a]

Sex	Mean	Std. Error	df	95% Confidence Interval	
				Lower Bound	Upper Bound
Female	6.050	.615	6	4.545	7.555
Male	4.881	.615	6	3.376	6.387

a. Dependent Variable: Y.

Pairwise Comparisons[b]

(I) Sex	(J) Sex	Mean Difference (I-J)	Std. Error	df	Sig.[a]	95% Confidence Interval for Difference[a]	
						Lower Bound	Upper Bound
Female	Male	1.169	.870	6	.228	-.960	3.298
Male	Female	-1.169	.870	6	.228	-3.298	.960

Based on estimated marginal means

a. Adjustment for multiple comparisons: Least Significant Difference (equivalent to no adjustments).

b. Dependent Variable: Y.